

In the Abstract:

Replace the Abstract as follows:

Various approaches for learning probabilistic relational models (PRM) having attribute uncertainty are described. In one approach, a parameter estimation task is provided by inputting a relational schema that specifies a set of classes. Attributes are associated with the classes, and the relational schema includes relationships between objects in different classes. A fully specified instance of the schema is provided in the form of a training database, and a structure learning task is performed to extract an entire PRM solely from the training database.

The invention comprises a method and apparatus for learning probabilistic models (PRM's) with attribute uncertainty. A PRM with attribute uncertainty defines a probability distribution over instantiations of a database. A learned PRM is useful for discovering interesting patterns and dependencies in the data. Unlike many existing techniques, the process is data driven rather than hypothesis driven. This makes the technique particularly well suited for exploratory data analysis. In addition, the invention comprises a method and apparatus for handling link uncertainty in PRM's. Link uncertainty is uncertainty over which entities are related in our domain. The invention comprises of two mechanisms for modeling link uncertainty: reference uncertainty and existence uncertainty. The invention includes learning algorithms for each form of link uncertainty. The third component of the invention is a technique for performing database selectivity estimation using probabilistic relational models. The invention provides a unified framework for the estimation of query result size for a broad class of queries involving both select and join operations. A single learned model can be used to efficiently estimate query result sizes for a wide collection of potential queries across multiple tables.